This summer has been an incredibly rewarding research journey: one that has pushed me to grow in confidence, problem-solving, and technical skills in ways I never anticipated. I began this project with a broad and somewhat ambitious goal: to investigate gender bias in generative AI systems, and to explore whether it could be reduced through fine-tuning with equity-focused datasets. But what I gained through the Laidlaw experience was far beyond a set of results. I learned how to adapt when things don't go as planned, how to ask for help and connect with others, and how to think like a researcher methodically, critically, and with clarity of purpose.

One of the most significant ways I've grown is in communication and confidence. Before this project, I was quite hesitant to reach out to others. I thought I needed to figure everything out on my own and asking for help might make me look unprepared or not knowledgeable enough. But this summer completely changed that mindset. As I moved deeper into the research process, I realized that connection and collaboration are not just helpful - they are essential. I found myself emailing and calling librarians at the university for access to databases I had never used before and data analysis help, reaching out to professors and graduate students in adjacent fields to ask questions about methodology and for expert evaluation of my prompt set, and cold-messaging other researchers online to discuss AI model behavior. I also learned how to communicate more clearly and confidently with my supervisor, Dr. Jieying Chen, who was incredibly generous with her time and expertise. I began to appreciate that people genuinely want to support undergraduate research. They enjoy sharing their knowledge and expertise with a new generation of researchers, pointing you toward helpful resources, and asking questions that challenge your assumptions. But that kind of help only becomes available when you take the initiative to reach out first.

Equally important was the growth I experienced in problem-solving and research design. When I compare my final methodology to the one outlined in my original proposal, they look dramatically different and that's a good thing. Initially, I aimed to compare models trained on biased vs. feminist datasets. However, I quickly realized that fine-tuning large-scale models wasn't feasible within the timeframe, and some models wouldn't even allow basic dataset integration. Rather than giving up, I

pivoted. I shifted my research toward a prompt-based diagnostic toolkit that could test for subtle gender bias across models. This process required me to rethink the entire structure of my project and learn from the start to develop a scale in scientific and reliable way. I learned how to define bias, how to identify it in story-based outputs, and how to analyze it systematically. That pivot also meant changing supervisors mid-project. Locating a new advisor, explaining my shift in direction, and building a new meeting structure with Dr. Jieying Chen was a valuable lesson in navigating academic transitions. It was the first time I truly understood that changing your plan isn't scary. In fact, choosing a plan that *actually works for you* is one of the most important lessons in doing real research.

Throughout the process, I encountered countless smaller roadblocks. Models refused to output numbers. Stories were too short or repetitive. I had to scope down my ambitions more than once, reminding myself that a "perfect study" isn't always possible, but a transparent, replicable one is. Each problem became a chance to troubleshoot, adapt, and test again. Despite these challenges, I was still able to present wholesome and reliable prompts for testing AI models' biases in text outputs, grounded in both theory and data.

Another major area of development was my growth in technical research tools. I used SPSS to run statistical analyses across three dimensions of bias, including t-tests and ANOVA tests to compare performance across models. I also learned MAXQDA, which I used to code over 250 AI-generated stories for subtle stereotype patterns. These software tools weren't even in my original plan, but I picked them up along the way out of necessity, and now I feel confident using them in future research. Along with that, I improved my skills in literature review, academic formatting, and structuring multi-phase research reports.

In conclusion, I didn't just learn about AI bias this summer,I learned how to design a study, manage transitions, and reach out to the people who can help bring your ideas to life. This experience has made me a more capable, confident, and curious researcher. I'm proud of the work I've done, but more than that, I'm grateful for the opportunity. The Laidlaw Scholars Programme has given me the space and

support to explore questions I've cared about for years. And in doing so, it's given me the foundation I need to pursue research that matters.